



Open Science Grid

VM/Cloud workshop OSG All-Hands meeting

March 8th 2010
Sebastien Goasguen

Welcome and Logistics

Ruth Says "Hi and thanks for Coming, I will be available all week to chat !"

Slides and Agenda are on Indico:

<http://indico.fnal.gov/conferenceDisplay.py?confId=3168>


If you need a chair, grab one next door.

On indico are EVO and phone coordinates for remote participation.

"Chez Leon" Dinner is at 6:30, inscriptions are closed.



Ruth Pordes
100 x 141 - 18k - jpg
opendsi.org



Why are we here ?

Learn about clouds and various technologies

Discuss, agree, disagree on what clouds means for the grid and OSG in particular.

Find common agreement on what cloud capabilities the VO want to see on OSG ?

Start putting cloud technologies in the OSG footprint, spark grass roots efforts...

Take the best of clouds, the best of grids and become a better OSG :)

Links




An EGEE comparative study: Grids and Clouds - Evolution or Revolution ? (Nov 2008)

<https://edms.cern.ch/file/925013/3/EGEE-Grid-Cloud.pdf>

Above the clouds: A Berkeley view of cloud computing (Feb 2009)

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf>

The cloud: lots of things...



- <http://contactdubai.com/tag/saas-software-or-storage-as-a-service>

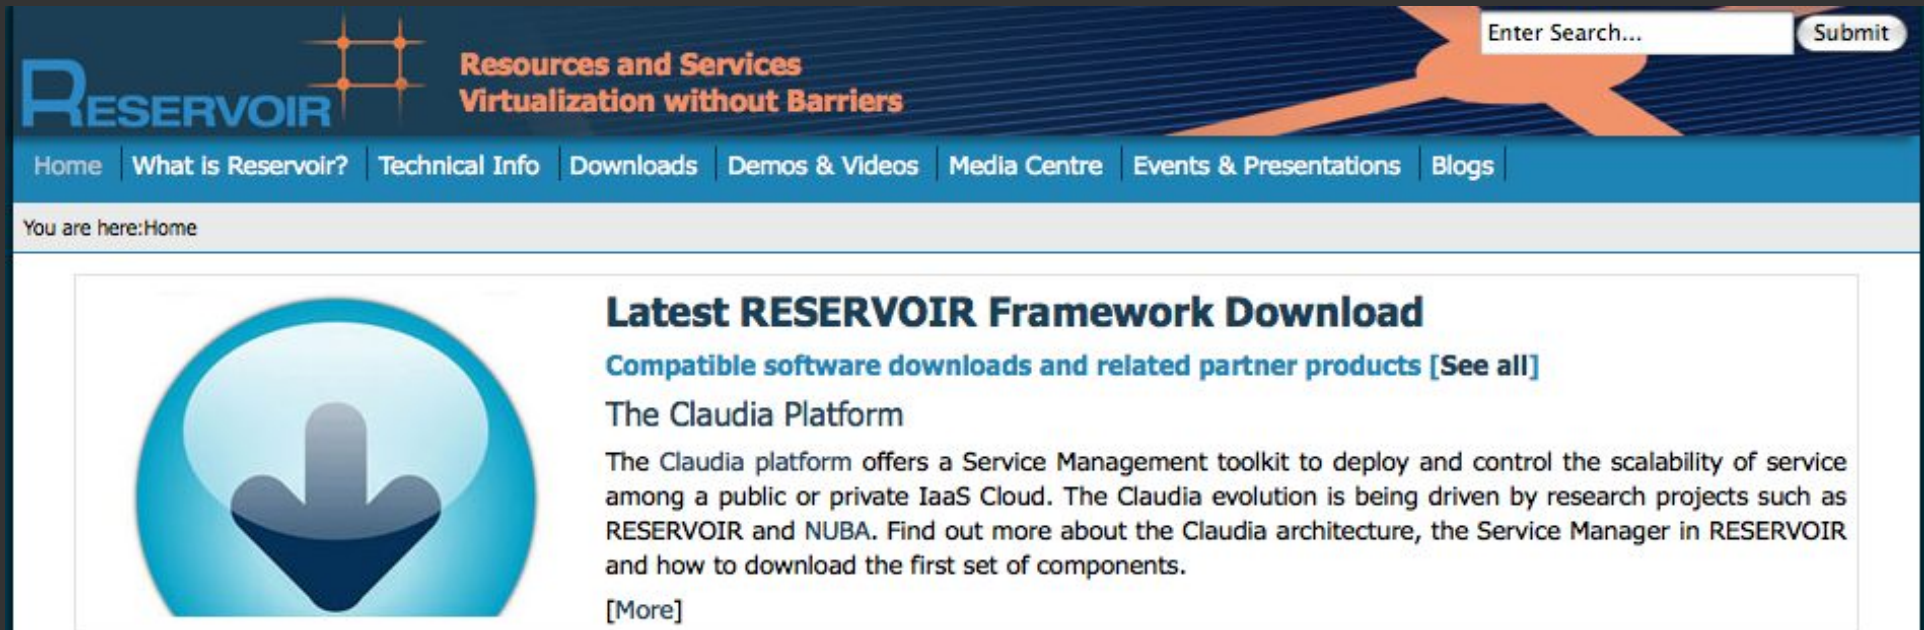
Agenda

- 13:30 Introduction and Open Nebula (20')** Sebastien Goasguen (*Clemson*)
- 13:50 CERN VM (30')** Predrag Buncic (*CERN*)
- 14:20 Status of the Nimbus Project (10')** Kate Keahey (*Argonne*)
- 14:30 Condor and VM's (10')** Jamie Frey (*University of Wisconsin, Madison*)
- 14:40 STAR's use of Virtualization and Clouds (10')** Levente Hajdu (*BNL*)
- 14:50 US ATLAS Plans and Thinking (10')** Yushu Yao (*LBNL*)
- 15:00 US CMS Plans and Thinking (10')** D. Evans and A. Melo (*Vanderbilt*)
- 15:10 Discussion of VO Commonalities (20')**
- 15:30 Break (30')** **16:00 Use of Future Grid (30')** Gregor Von Laszewski (*Indiana University*)
- 16:30 Magellan Project (30')** Susan Coghlan (*ANL*)
- 17:00 VO's Working Together (1h00')**
- Discussion and agreement on action items for OSG and VOs

This talk:

- Opennebula
- CERN Ixcloud effort
- Clemson cloud

Reservoir and Opennebula



The screenshot shows the RESERVOIR website homepage. The header features the RESERVOIR logo on the left, followed by the tagline "Resources and Services Virtualization without Barriers". On the right, there is a search bar with the placeholder text "Enter Search..." and a "Submit" button. Below the header is a navigation menu with links: Home, What Is Reservoir?, Technical Info, Downloads, Demos & Videos, Media Centre, Events & Presentations, and Blogs. A breadcrumb trail below the menu reads "You are here: Home". The main content area features a large blue circular icon with a white downward arrow. To the right of the icon, the text reads: "Latest RESERVOIR Framework Download", "Compatible software downloads and related partner products [See all]", "The Claudia Platform", and a paragraph describing the Claudia platform as a Service Management toolkit for IaaS Cloud scalability, driven by research projects like RESERVOIR and NUBA. A "[More]" link is at the bottom of the text block.

RESERVOIR Resources and Services
Virtualization without Barriers

Enter Search... Submit

Home | What Is Reservoir? | Technical Info | Downloads | Demos & Videos | Media Centre | Events & Presentations | Blogs

You are here: Home

Latest RESERVOIR Framework Download

[Compatible software downloads and related partner products \[See all\]](#)

The Claudia Platform

The Claudia platform offers a Service Management toolkit to deploy and control the scalability of service among a public or private IaaS Cloud. The Claudia evolution is being driven by research projects such as RESERVOIR and NUBA. Find out more about the Claudia architecture, the Service Manager in RESERVOIR and how to download the first set of components.

[\[More\]](#)

European Commission Project (17M Euros)
MoU between EGEE and Reservoir
Opennebula is one of the projects under Reservoir

Opennebula

OpenNebula.org

The Open Source Toolkit for Cloud Computing

- Home
- About
- Documentation
- Software
- Support
- Community
- Cloud
- Dev

The Leading and Most Advanced Solution for Cloud Computing

OpenNebula is the standard-based open-source toolkit to build private, public and hybrid clouds. The toolkit provides flexible architecture, interfaces and components that fits into any existing data center. Key features in OpenNebula v1.4 are:

- support for the **Xen**, **KVM** and **VMware** virtualization platforms.
- Access to **Amazon EC2** and **ElasticHosts** clouds.
- **libvirt**, **EC2 Query API** and **OGC OCCI** interfaces.

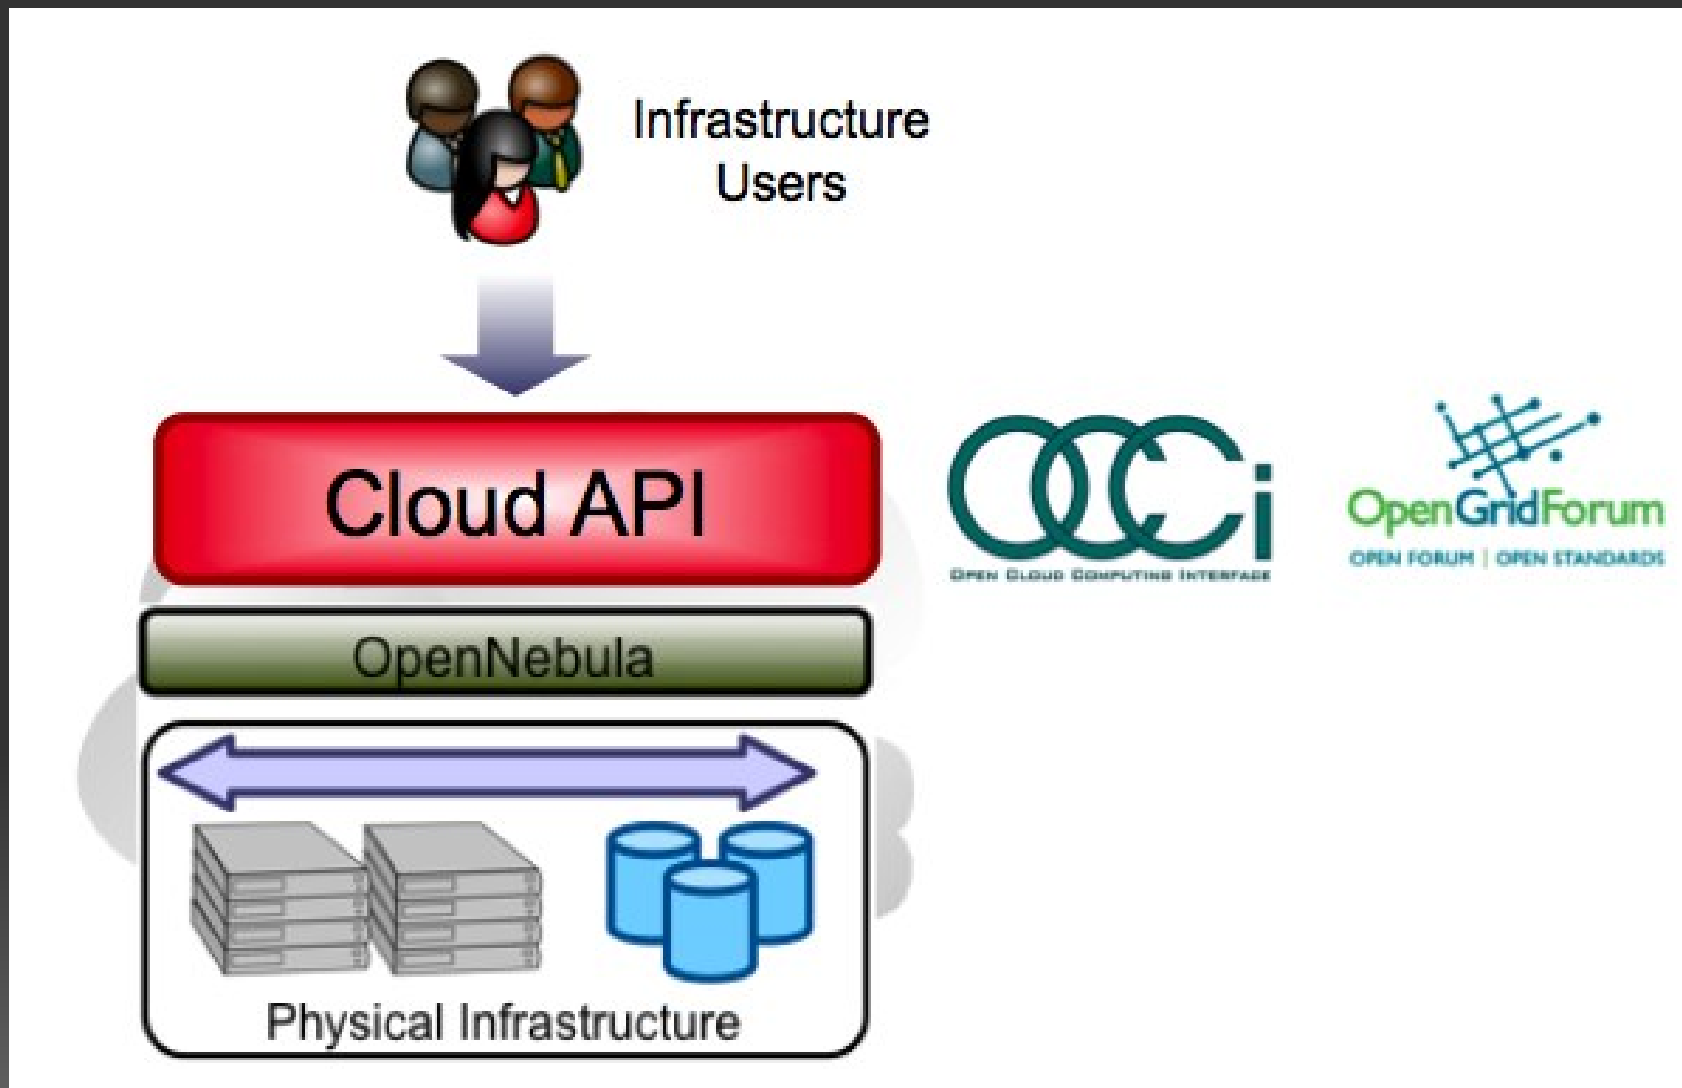
The diagram consists of four blue cloud shapes. A central cloud is labeled 'OpenNebula'. To its top-left is a cloud labeled 'Plugins', to its top-right is a cloud labeled 'Tools', and to its right is a cloud labeled 'Extensions'.

Announcements

- [New Web Site for OpenNebula.org](#) 2010/02/24
- [OpenNebula 1.4 Released](#) 2009/12/16
- [OpenNebula Cloud](#) 2009/11/18 Announcement
- [OpenNebula 1.4 RC Released!](#) 2009/11/18

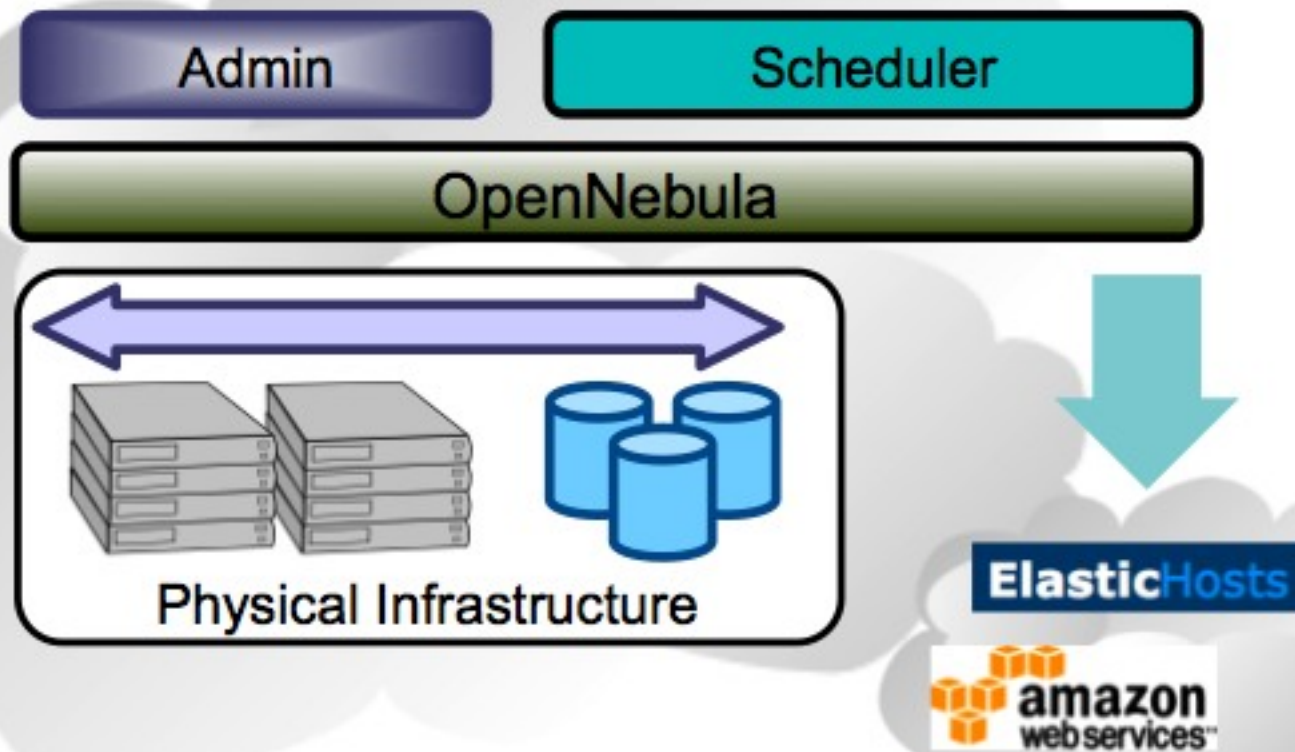
Out of University Complutense of Madrid
Ignacio de Llorente and Ruben Montero with team (Tino, Javi ...)
Very active projects (mailing lists, twitter, fb...)

Opennebula (Private, Public and Hybrid clouds)



Opennebula

Mostly Ruby + shell scripts + lots of ssh



one cli ...and also an api

```
-bash-3.2$ onehost list
```

HID	NAME	RVM	TCPU	FCPU
3	lxbrl2302	0	800	798
4	lxbrl2303	0	800	798
5	lxbrl2304	0	800	799
6	lxbrl2305	0	800	799
8	lxbrl2307	0	800	799
9	lxbrl2308	0	800	798
10	lxbrl2309	0	800	797
11	lxbrl2310	0	800	799
12	lxbrl2311	0		0
13	lxbrl2312	0	800	799
14	lxbrl2313	0	800	799
15	lxbrl2314	0	800	799
16	lxbrl2315	0	800	799
17	lxbrl2316	0	800	798
18	lxbrl2317	0	800	799

```
-bash-3.2$
```

```
-bash-3.2$ onevnet list
```

NID	NAME	TYPE	BRIDGE
1	Public	Fixed	eth0

```
-bash-3.2$ onevnet show 1
```

NID : 1
UID : 0
Network Name : Public
Type : Fixed
Bridge : eth0

```
....: Template :....
```

```
BRIDGE=eth0
```

```
LEASES=IP=00:16:3E:00:02:3B,MAC=10.12.3.30
```

```
LEASES=IP=00:16:3E:00:02:3C,MAC=10.12.3.31
```

```
LEASES=IP=00:16:3E:00:02:3D,MAC=10.12.3.32
```

```
LEASES=IP=00:16:3E:00:02:3E,MAC=10.12.3.33
```

```
LEASES=IP=00:16:3E:00:02:3F,MAC=10.12.3.34
```

```
LEASES=IP=00:16:3E:00:02:40,MAC=10.12.3.35
```

```
LEASES=IP=00:16:3E:00:02:41,MAC=10.12.3.36
```

```
LEASES=IP=00:16:3E:00:02:42,MAC=10.12.3.37
```

```
LEASES=IP=00:16:3E:00:02:4A,MAC=10.12.3.45
```

```
LEASES=IP=00:16:3E:00:02:49,MAC=10.12.3.44
```

```
LEASES=IP=00:16:3E:00:02:48,MAC=10.12.3.43
```

```
LEASES=IP=00:16:3E:00:02:47,MAC=10.12.3.42
```

```
LEASES=IP=00:16:3E:00:02:46,MAC=10.12.3.41
```

```
LEASES=IP=00:16:3E:00:02:45,MAC=10.12.3.40
```

```
LEASES=IP=00:16:3E:00:02:44,MAC=10.12.3.39
```

```
LEASES=IP=00:16:3E:00:02:43,MAC=10.12.3.38
```

```
NAME=Public
```

```
TYPE=FIXED
```

```
....: Leases :....
```

```
-bash-3.2$
```

Ixcloud at CERN

CERN IT is investigating how to support VO based images by virtualizing its computing farm.

First steps are to operate hypervisors at a large scale (~2,000) and run worker nodes as virtual machines.

Ixcloud has been created with a few hosts (10) to develop the operational processes and tests various pieces of the puzzle (image management, image distribution etc...)

Ixcloud is testing both opennebula and Platform ISF solutions.

IT-pes group: Ulrich Schwickerath, Ewan Roche, Belmiro Moreira, Sebastien Goasguen, see Hepix Fall 2009 for details

Main ideas

Golden nodes managed by Quattor. VMs are not.

Golden nodes snapshotted and distributed (?) to the hypervisors


LVM snapshots used to speed up VM boot time on hypervisors

VM worker node can be set to drain, once no jobs are running the VM shuts down via a call to opennebula.


LSF queue is inspected to decide what the VM pool should be made of (SLC4, SLC5, CERNVM ..etc) ? Opennebula is used to create that VM pool in steady state.

Mid/Longer term: Allow VOs to use their own image, potentially upload their image a la cloud.


CERN early results



CERN early results



CERN early results



Clemson setup

Virtual Organization Cluster

- > Participating VOs: Give an image
- > Non-Participating VOs: Run in a std image without knowing
 - >transparent to users, they use the same job submission techniques than on the grid

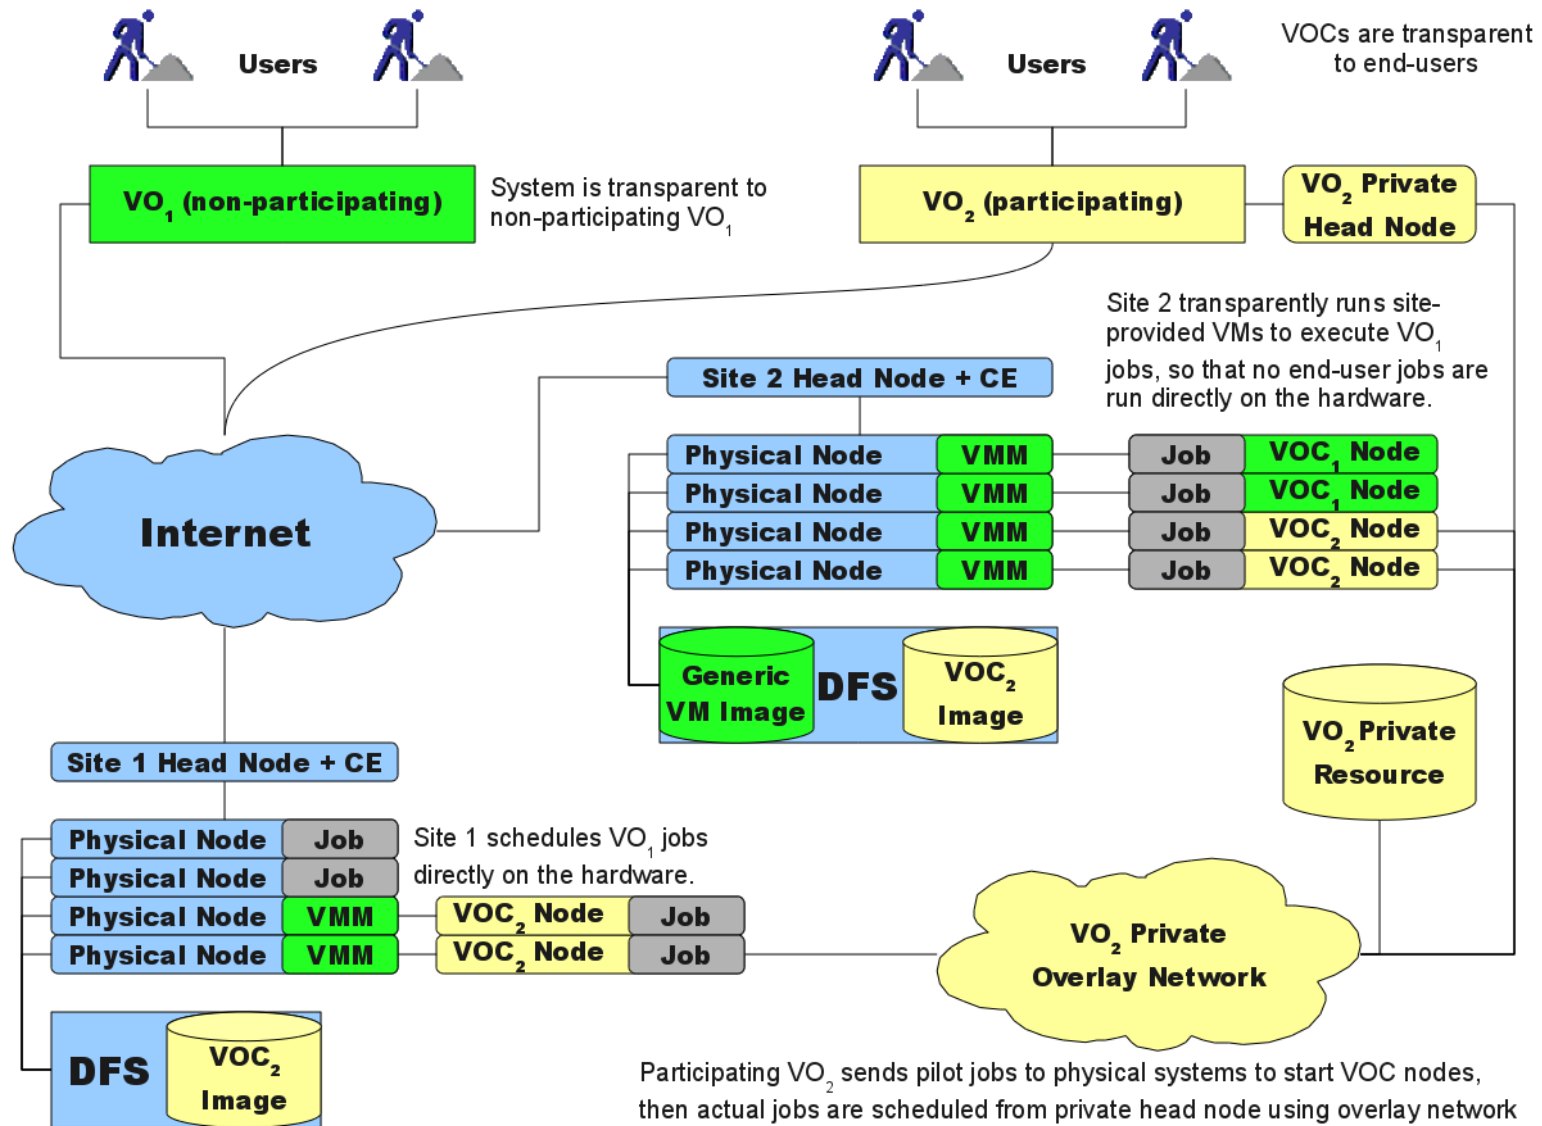
A regular cluster on OSG that uses KVM

VMs are booted according to the workload submitted to the cluster (The pool size will track the job load)


M. A. Murphy, L. Abraham, M. Fenn and S. Goasguen

*“Autonomic Clouds on the Grid” Journal of Grid Computing
Volume 8, Number 1 (March 2010), pages 1-18.*

Vision for Virtual Organization Clusters




Clemson OSG cloud results



Engage (short test)

Clemson OSG cloud (2-1/2 months trace)



Scaling this to production cluster

Worked with IT to deploy KVM on a 8,000 core cluster at Clemson.

Run virtual machines like any regular job on the cluster. Use a PBS script to get nodes and start the KVM VMs on each node (8 VMs per node, one per core).

Scaled to ~1,000 VMs started. Suprisingly we did this using a single image in /home and using the KVM snapshot mode that create qcow images on the hosts. **See e-demo tuesday**

Currently testing new job scheduling framework to forget about the NAT networking :)

Thank You

Questions ?

sebgoa@clemson.edu

or

mfenn@clemson.edu (Needs a job in the next 3 days)

or

mamurph@clemson.edu

Otherwise let's stay on time...